LEVEL III

12

INTERACTIVE AIDS FOR CARTOGRAPHY AND PHOTO INTERPRETATION

A064088

AD A073452

Semiannual Technical Report
Covering the period October 12, 1978 to May 11, 1979

SRI Project 5300
Contract No. DAAG29-76-C-0057

ARPA Order No. 2894-5
Program Element Code 61101E

Contract Amount: $1,373,527
Effective Date: May 12, 1976
Expiration Date: October 9, 1979

By: Martin A. Fischler, Senior Computer Scientist
Principal Investigator
(415) 326-6200, Ext. 5106

and

G. J. Agin, L. H. Quam
H. G. Barrow, J. M. Tenenbaum
R. C. Bolles, H. C. Wolf

Prepared for:

Director
Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, Virginia 22209

Attention: Program Management

Approved for public release; distribution unlimited.

SRI International
333 Ravenswood Avenue
Menlo Park, California 94025
(415) 326-6200
Cable: SRI INTL MPK
TWX: 910-373-1246

SRI International

SRI International

6

INTERACTIVE AIDS FOR CARTOGRAPHY AND PHOTO INTERPRETATION

9 Semiannual Technical Report, 12 Oct 78 - 11 May 79,
Covering the period October 12, 1978 to May 11, 1979

SRI Project 5300
Contract No. DAAG29-76-C-0057, ARPA Order-2894-5
15

ARPA Order No. 2894-5
Program Element Code 61101E

Contract Amount: $1,373,527
Effective Date: May 12, 1976
Expiration Date: October 9, 1979

11 11 May 79

12 31p.

10

By: Martin A. Fischler, ~~Senior Computer Scientist~~
~~Principal Investigator~~
~~(415) 326-6200, Ext. 5106~~

and

G. J. Agin, L. H. Quam
H. G. Barrow, J. M. Tenenbaum
R. C. Bolles, H. C. Wolf

Prepared for:

Director
Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, Virginia 22209

Attention: Program Management

Approved:

Peter E. Hart, Director
Artificial Intelligence Center

David H. Brandin, Executive Director
Computer Science and Technology Division

410 281

JOB

ABSTRACT

The central scientific goal of the ARPA Image-Understanding Project research program at SRI International is to investigate and develop ways in which diverse sources of knowledge may be brought to bear on the problem of interpreting images. The research is concerned with specific problems that arise in processing aerial photographs for such military applications as cartography, intelligence, weapon guidance, and targeting. A key concept is the use of a generalized digital map to guide the process of image analysis.

In the present phase of our program, the primary focus is on developing a "road expert," whose purpose is to monitor and interpret road events in aerial imagery. The objectives, methodology, and current status of our research are described in this report. Particular technical topics include data base construction and shadow and anomaly analysis.

CONTENTS

ILLUSTRATIONS

ACKNOWLEDGMENTS

The work reported here was performed by G.J. Agin, H.G. Barrow, R.C. Bolles, M.A. Fischler, L. Quam, J.M. Tenenbaum, and H.C. Wolf.

The research was supported by the Defense Advanced Research Projects Agency of the Department of Defense and was monitored by the U.S. Army Research Office under Contract No. DAAG29-76-C-0057.

I DETECTING AND INTERPRETING ROAD EVENTS IN AERIAL IMAGERY

A. Introduction

The central scientific goal of the ARPA Image-Understanding Project research program at SRI International is to investigate and develop ways in which diverse sources of knowledge may be brought to bear on the problem of interpreting images. The research is concerned with specific problems that arise in processing aerial photographs for military applications such as cartography, intelligence, weapon guidance, and targeting. A key concept is the use of a generalized digital map to guide the process of image analysis.

In the present phase of our effort, the primary focus is on developing a "road expert," a computer program whose purpose is to monitor and interpret road events in aerial imagery.

Our significant accomplishments include:

1) The introduction and exploitation of two major paradigms:

   a) Map-Guided Image Interpretation--Establishing a projective correspondence between a symbolic data base and an image, and using the data base to guide and constrain the interpretation of the image.

   b) Perceptual Reasoning--Modeling the information sources and image operators so that selection of analysis techniques, location of search areas in the image, sequencing of information acquisition, and the way in which perceived and a priori information are combined into a final interpretation are matched to scene content and viewing conditions.

2) The design and implementation of the SRI Road Expert--a framework for understanding the requirements for achieving human-like performance in the analysis of aerial imagery.

The task of road monitoring provides the context for this investigation. Our work has concentrated on three major subtopics: establishing a correspondence between an image and an existing map data base; detecting and delineating the visible roads; identifying the objects appearing on and along the road surfaces. Our specific objectives, approach, and progress are described below.

B. Objective

The primary objective of this research is to build a computer system that "understands" the nature of roads and road events. It should be capable of performing such tasks as:

(1) Finding roads in aerial imagery.

(2) Distinguishing vehicles on roads from shadows, signposts, road markings, etc.

(3) Comparing multiple images and symbolic information pertaining to the same road segment, and deciding whether significant changes have occurred.

The system should be capable of performing the above tasks even when the roads are partially occluded by clouds or terrain features, are viewed from arbitrary angles and distances, or pass through a variety of terrain.

C. Approach

To achieve the above capabilities, we are developing two "expert" subsystems: the "Road Expert" and the "Vehicle Expert." The Road Expert knows mainly about roads, how to find them in imagery, and what things belong on them. It works at low-to-intermediate resolution (e.g., 1-20 ft. of ground distance per image pixel) and has the ability to distinguish vehicles from other road detail. The Vehicle Expert works on higher-resolution imagery and can identify vehicles as to type. We are concentrating our efforts on the Road Expert and therefore will limit most of our discussion here to this component of our system.

The major tasks performed automatically by the Road Expert are:

(1) Image/map correspondence--Placing a newly acquired image into geographic correspondence with the map data base.

(2) Road tracking--Precisely marking the center line of selected visible sections of road in the image.

(3) Anomaly analysis--Locating and analyzing anomalous objects on, and adjacent to, the road surface; identifying potential vehicles.

The image/map correspondence task is accomplished by locating roads and road features as landmarks; correspondence is performed at resolutions as coarse as 20 ft./pixel, so that a reasonably wide field of view (10 to 100 sq. mi.) can be processed at one time. It is nominally assumed that the initial combinations of uncertainties as to the estimates for the camera parameters imply uncertainties on the ground of approximately +/- 200 ft. in X and Y. The correspondence procedure works iteratively to refine the camera parameters. A typical goal is to reduce the implied uncertainties on the ground to about +/- 2 ft. in X and Y.

After the image is placed into correspondence with our map data base, one or more of the visible road sections are selected for monitoring. The road centerline and lane boundaries are found to an accuracy of one to two pixels in imagery with a resolution of 1 to 3 ft./pixel.

Given the precise road locations in the image, anomalous objects are detected by scanning on and along the road pavement. These anomalous objects are then identified as to type (e.g., vehicle, shadow, road surface marking, signpost, etc.).

The above tasks are supported by information about the road's condition and general structure from a symbolic data base. For example, if prior photographic coverage of the area being analyzed is available, the problem of anomaly classification can be simplified by determining whether a similarly shaped anomaly can be found in the same general location over some prolonged period. Additional examples of how data base knowledge and stored models can aid in the analysis process

include: using the time of day in discriminating shadows from objects of interest; utilizing the general shape and width of the road (obtained from a map) as an aid in road tracking; providing relevant information on the anticipated size, shape, and road orientation of potential vehicles.

A central theme of this effort is to consider road monitoring as a knowledge domain. In particular, we are addressing ourselves to the question of how a priori knowledge can be directly invoked by the image-analysis modules (what type of knowledge, how it should be represented, and what mechanisms there are for its use). To achieve our goal of building a very-high-performance system, we are developing explicit models of the image structures we are dealing with and, additionally, models of the decision procedures embedded in the image-processing algorithms, so that the algorithms can evaluate their own performance. Finally, we are planning an overall control structure that will be concerned with the problems of coordinating analysis across a spectrum of resolution levels, as well as with those of integrating multisource information.

D. Progress

Our work to date has provided the capabilities necessary to assemble an integrated Road Expert demonstration system, and we are currently planning to have such a system operational by October 1979. This system will allow a user to submit new photographs from a previously "instantiated" site for automatic analysis, in which image scanning, image-to-data base correspondence, road marking, and anomaly analysis will be performed "on line".

The demonstration system will also permit both interactive instantiation of a new site and selected analysis functions (such as road tracking) on photographs for which there is no data base support.

We have previously described [2, 3] our approach to the correspondence and road marking tasks; work continues in these two areas, not only to achieve higher performance, but also to generalize

the techniques to a wider class of domains. A more detailed description of this continuing work will be deferred until a later time.

In the following two subsections we shall describe recent progress in dealing with the problem of vehicle detection and anomaly analysis; we shall also discuss our plans for on-line site instantiation.

1. Progress in Anomaly Classification

We now have a program that will analyze the anomalies detected by the correlation road tracker [3] and decide whether or not they result from vehicles. If an anomaly is judged to be a vehicle, then the program will provide a limited amount of classification as to vehicle type. If the anomaly is judged to be something other than a vehicle, the program provides the most likely interpretation of what it is.

The correlation road tracker has been modified to produce, in addition to the road track, an image array containing the difference between the actual brightness in the original image and the brightness predicted from the road model (originally this additional output was in the form of a binary anomaly mask). The value of this "difference image" is twofold: it can be thresholded to decide what is or is not anomalous, and the image with the road profile excluded is useful for analyzing shadows and road discolorations.

It is obvious that an understanding of shadows is crucial in making sense out of road scenes. Aerial scenes are often photographed in direct sunlight, and vehicles on the road cause anomalies that include the vehicle plus its shadow. Large objects off the road, such as signs, trees, and utility poles cast shadows that are noticed by the anomaly detector. In addition, the shadows can give valuable clues as to the size and shape of the objects casting them.

We employ three basic techniques for identifying shadows. A brightness model allows us to identify shadows by the absolute brightness of pixels in the difference image. A predictive model allows us to identify the portion of an anomaly most likely to be shadow when

we know the position of the sun and the height of the object casting the shadow. Finally, a projective model, which tries to detect the two long parallel sides of a vehicle, can locate the dividing line between a vehicle and its shadow.

A number of "expert subroutines" examine each anomaly. The vehicle expert subroutine exploits the basically rectangular shape of vehicles when viewed from above. Anomalies that are clearly the wrong size are eliminated at the outset. Projecting the average brightness and average gradient magnitude upon a baseline perpendicular to the presumed direction of vehicle travel enables location of the shadow and establishment of a nominal width for the vehicle. Height can usually be estimated from the shadow, and length is inferred from the size of the total anomaly (allowing for a shadow fore or aft).

Two other anomaly experts, the tree-shadow expert and the road marking expert, provide alternate explanations for anomalies not identified as vehicles. To qualify as a tree shadow (or the shadow of some other object off the road) an anomaly must have the appropriate average brightness, a low variance in brightness, and touch the side of the road at the side nearer the sun. Road markings (as a rule, painted arrows or speed limit numerals) are usually brighter than the road surface, have low brightness variance, and are quite limited in extent.

A detailed discussion of the above material is contained in Section II of this report.

2. The Road Data Base and its Compilation

This subsection describes the present state of implementation of the road data base and plans for the October 1979 demonstration involving on-line site instantiation.

The purpose of the road data base is to enable the Road Expert to find known roads in new images accurately and reliably, trace their paths, and locate anomalies that might be potential vehicles on the roads. The data base also contains information to help distinguish

vehicles from such permanent road features as signs and their shadows, and painted markings on the road surface.

The current road data base contains both geometric and photometric information. The geometric part of the road data base was generated by a variety of means, depending on the level of detail and accuracy desired. The coarsest level of data representation was generated by specifying approximate world location, direction, and width of road segments, either by typing in numerical information or by tracing the road in a low-resolution (USGS 7.5 minute series) map of the area. The most accurate geometric information was entered into the data base both by typing in precise numerical data and by manually tracing portions of "as built" survey plans of the road obtained from the California Department of Transportation.

Photometric information associated with a road segment is inserted into the data base by using the correlation road tracker; as images of a geographic site are interpreted by the road tracker, road photometry models are automatically entered. Spatially fixed landmarks, such as painted road-surface markings, are (at present) manually specified; and a corresponding rectangular image patch is entered into the data base.

The data base is currently implemented by means of SAIL record structures that conveniently provide graph structures, lists, numeric arrays, etc. A general-purpose record structure I/O package communicates these structures between SAIL programs and disk files. We recognize the eventual need to develop a file representation that can be communicated to LISP programs.

We intend to include examples of data base construction as a part of the Road Expert demonstration and are working toward a scenario of the following type. An image of a site will be scanned and digitized at approximately 1-3 ft. per pixel resolution; a photo interpreter will then indicate the approximate locations of primary road segments in the image, using a track ball. The automatic road-tracker program will be invoked to accurately trace the roads, generate cross-section photometry

models, and detect anomalies that might be permanent surface markings. The anomaly analysis techniques described in the preceding subsection (and in Section II) will specify which anomalies are to be included as point features in the data base. The photo interpreter will then review and edit the results.

Since a single image will not provide terrain elevation information, we are hoping to proceed as follows. After one image of a stereo pair has been analyzed as described above, the second image of the pair will be scanned and digitized. The second image will be used to determine relative elevations of road points by parallax measurements made on road surface features or nearby image areas that can be aligned by cross-correlation. Real world x,y,z will be determined from knowing the world location of a few recognizable landmarks in the images.

E. Comments

We see the military relevance of our work extending well beyond the specific road-monitoring scenario presented above. In particular, a Road Expert can be applied to such problems as:

(1) Intelligence--Monitoring roads for movement of military forces

(2) Weapon guidance--Use of roads as landmarks for "map-matching" systems

(3) Targeting--Detection of vehicles for interdiction of road traffic

(4) Cartography--Compilation and updating of maps with respect to roads and other linear features (especially those concerned with transportation), such as airport runways, railroads, rivers, etc.

In accordance with our generalized view of the applicability of the Road Expert and the knowledge-based, image-analysis techniques we are developing, we are attempting to achieve a level of performance and understanding in each functional task far exceeding that required for dealing with the road-monitoring scenario alone.

The remainder of this report presents a detailed discussion of our current work on the problem of detecting and analyzing objects appearing on and along the roads being monitored.

II KNOWLEDGE-BASED DETECTION AND CLASSIFICATION OF VEHICLES AND OTHER OBJECTS IN AERIAL ROAD IMAGES

A. Introduction

This section describes an approach to finding and identifying vehicles in aerial images, using diverse sources of knowledge. The following scenario provides a context for this work. Given a digital aerial image and a data base, the problem is to detect vehicles on the road and to classify them as to vehicle type. The image should have sufficient spatial resolution to allow recognition (about one ft. per pixel, minimum). Figure 1 shows a typical image of an area containing a freeway interchange.

The data base contains information about some limited geographical area of interest. As a minimum, it should have the locations of known roads in the area. Other relevant information could include (but not be limited to):

* Road width
* Brightness profiles across the road
* Terrain information
* Buildings, railroads, and other cultural features
* Intersections, overpasses, and access roads
* Signs and permanent road markings
* Previous photo coverage of the area, in digital form.

Figure 1 An Aerial Road Image

A calibration procedure [6] establishes correspondence between image coordinates and geographic coordinates, allowing us to convert quickly back and forth between coordinates in the data base and pixel locations in the image. A road tracker [3] uses the road location predicted by the data base to trace the road centerline and boundaries by correlating successive profiles perpendicular to the road direction. Areas where the image diverges from the expected road profile are identified as "anomalies." These areas are passed to the classification routines for further scrutiny.

Many different conditions could give rise to an anomaly. Vehicles usually show up this way, but so do the shadows of objects off the road (trees, buildings, signs, utility poles), overhanging trees, painted markings on the road, and changes or irregularities in the road surface

(such as tar patches). There are also some less frequent situations with which a practical system ought to deal, such as road construction, floods, bomb craters, smoke, and dust clouds. The classifier must first decide if the anomaly arises from a vehicle or from some other cause. Then it can classify the vehicle type.

Although the scenario assumes some rather specific resources and goals, this knowledge-based approach is generally applicable to a wide range of object recognition tasks in cartography and photo interpretation.

B. Sources of Information

A wide variety of information can be helpful for detecting and classifying vehicles. We can identify three kinds of knowledge relevant to this problem: about the problem domain (generic knowledge), about the site (the data base), and about a particular place and time (information associated with the image).

Generic knowledge includes information that can be deduced from functional descriptions. A road is a narrow, linear region upon which vehicles may travel. The road is usually continuous in the image--if it appears discontinuous it may be that there are obstructions, or there may be shadows or discolorations on the road surface. Roads have minimal variation in the direction of travel but may have considerable variation in the perpendicular direction, because of the different compositions of roadbed, shoulders, and an expected pattern of oil stains in the center of each lane. We have some idea of the expected shapes of vehicles viewed from different angles, and an expectation that they probably will be aligned parallel to the road direction. Our illumination models take into account the physics and geometry of shadows, and we can sometimes use shadows to draw inferences about objects. We know the usual places where road signs, utility poles, and painted road markings are located. All the foregoing can be used to make sense out of a road scene.

The data base is a useful source of information. Its principal use is to predict the approximate road centerline, so that the road-tracking subroutines can operate. But other kinds of information can be brought into play. Terrain information can be used to refine position estimates when the viewing angle is not vertical and to predict shadows better if the ground slopes. Classifying shadows of objects off the road is very much simplified when it is known what objects are likely to cast shadows. Ambiguous anomalies in the image can sometimes be distinguished if a picture can be compared with a previous one or, better yet, if the data base states what anomalies were found in previous images and how they were classified. Intelligence reports and expected traffic conditions can help the program decide what to look for or what strategies to use.

The greatest single source of data is the image itself. It is easy to overlook some information that is associated with the image but may not be in the actual raster. For example, it is usually possible to ascertain (at least approximately) the altitude, position, and heading of the aircraft from which the image was taken. Scaling parameters, view angles, and compass headings can be derived by calibration. If the time and date of the picture are known, the sun position can be calculated--but even without these data the sun position usually can be estimated from shadows.

In short, detection and classification of vehicles are not based solely on what is in the image. In the following sections, we detail some of the ways we use the available information.

C. Use of the Correlation Road Tracker

We depend on the correlation road tracker designed by Quam [3] to isolate anomalies in images of roads. These are regions where attention should be focused.

The road tracker is based on the assumption that variations in road surface materials, centerlines, and intralane wear patterns correspond linearly to the road itself. Vehicles and other anomalies, however,

stand out in sharp contrast to the pattern of the road. Detecting these anomalies is important to the operation of the road tracker. Where substantial disagreement occurs between successive profiles, the corresponding pixels are marked as anomalies, so that these points can be eliminated from the correlation calculations. If the anomalies were not so masked, they would perturb the location of the correlation peak and introduce errors.

Figure 2a shows a representative excerpt from the area covered by the image of Figure 1. The road tracker is initiated by specifying a single profile approximately perpendicular to the road direction and centered on it. This initial baseline is now selected manually, but facilities exist for using the data base to draw the baseline automatically.

The road tracker produces several forms of output. As indicated by Quam [3], the program can produce a point list describing the track of the road center, as well as a binary image of all points in the road that are anomalous. But for vehicle identification another form of output has been added. The road reflectance model may be subtracted from each pixel considered, resulting in a difference image that has been normalized to remove the road profile. Figure 2b shows the baseline, the road center, and anomalies detected. Figure 2c shows the difference image. The difference image may be converted to a binary anomaly image by thresholding.

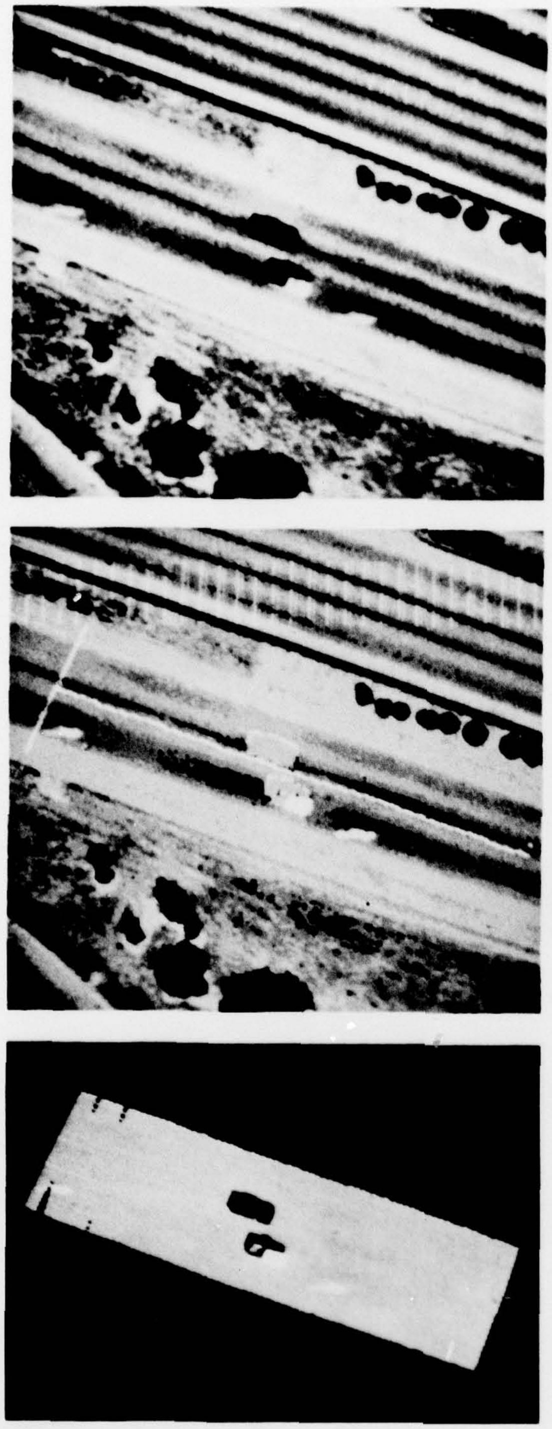

(a) ROAD SCENE

(b) BASELINE, CENTERLINE, AND ANOMALIES

(c) DIFFERENCE IMAGE

Figure 2 Operation of the Road Tracker

In the difference image, shadows tend to have a relatively uniform intensity, even though the road reflectance profile varies considerably. If we adopt the simplifying assumptions that any object casting a shadow may be approximated by a half plane of infinite extent that hides all but a fixed proportion of the sky, and if we neglect reflected illumination from nearby objects, then the ratio of intensities across the shadow edge should not depend on the reflectivity of the underlying surface. When the original image is digitized on a logarithmic brightness scale, this constant ratio becomes a constant intensity in the difference image. Because the assumptions are approximate at best, the constant-difference test is almost never exact. Nonetheless, by subtracting the road profile from the image, we can expect the intensity of shadows to be more uniform in the difference image than in the original one.

On the other hand, when anomalies are caused by vehicles, subtracting the road profile will cause its inverse to be superimposed on the anomaly. Figures 3a and b show an original image and a difference image (from another road site) that demonstrate these peculiarities. Both kinds of image are useful in classifying anomalies.

As the road tracker proceeds, it constantly keeps track of the average correlation between successive road profiles at their optimum locations. This correlation value, a useful estimate of noise in the picture, is made available to succeeding classification stages.

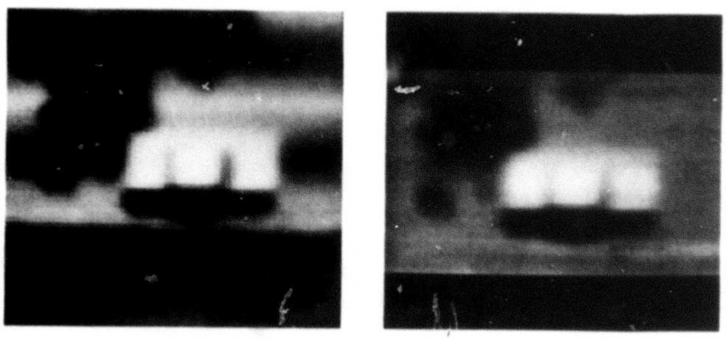

(a) ORIGINAL IMAGE (b) DIFFERENCE IMAGE

Figure 3 Original and Difference Image

D. Shadows

An understanding of shadows is crucial to making sense out of high-resolution aerial images. The scene is always out-of-doors and usually illuminated by direct sunlight, which produces deep, dark shadows. Frequently shadows are the most prominent visual feature of an image.

For vehicle classification, many of the anomalies the classifier is called on to consider are the shadows of objects off the road, such as trees, signs, or utility poles. All vehicles cast shadows, and, unless the boundary between the vehicle and its shadow can be determined, classification on the basis of shape is hopeless. Furthermore the existence or nonexistence of a shadow can aid in deciding whether or not a given anomaly is a vehicle. The size and shape of the shadow can give valuable clues as to the height of the vehicle and its profile. As a dramatic demonstration of this, consider the vehicle shown in Figure 4. Because its reflectance is almost the same as that of the road, the vehicle might have gone unnoticed, were it not for the shadow. But the shadow not only gives away its position; it tells us the vehicle is probably a Volkswagen "beetle."

Figure 4 Vehicle with Shadow

We have a number of techniques at our disposal for identifying shadows. The simplest is based on the brightness model. The technique is simply to search for all pixels in the image whose intensity is in the range of values expected for shadows. This works somewhat better in the difference image than in the original, because the effects of variation in the road surface are reduced. Figure 5 shows the central portion of the area analyzed in Figure 3, which we shall use to illustrate shadow-finding techniques. Figure 6 shows the shadows extracted from Figure 5b by this method.

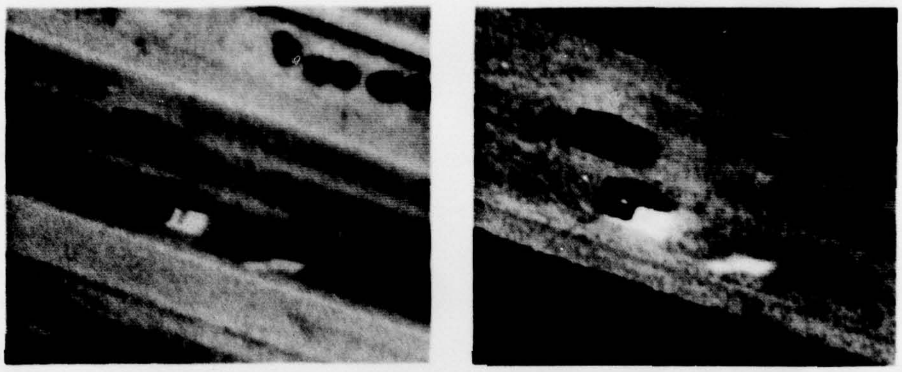

(a) ORIGINAL IMAGE (b) DIFFERENCE IMAGE

Figure 5 Original and Difference Pictures

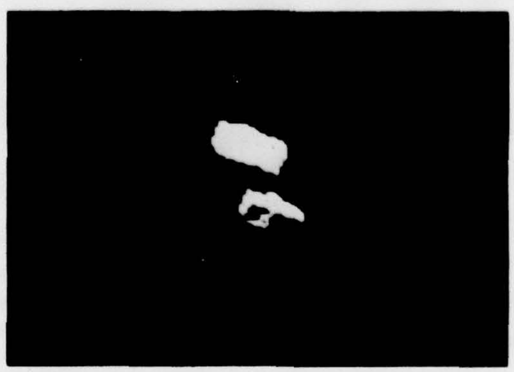

Figure 6 Shadows Found by Brightness Criterion

In our work so far, the expected range of shadow intensities has been inferred from the statistics of areas manually indicated as shadows. It should be possible in principle to automate this procedure--for example, by using the using the data base to predict or find known shadows. Alternatively, it seems likely that a formula can be derived that will give the expected distribution based on calibration of photometry.

In situations in which the correlation road tracker is not applicable, shadows located by the brightness model might indicate areas of the picture that merit scrutiny.

Another device, based upon a predictive model, depends on knowing the sun's angle. The shadow of any raised object is always on the side away from the sun; and, if the height of the object is known, the length of the shadow can be predicted. Figure 7 shows the areas identified as shadow from the image of Figure 5b by thresholding the difference image to locate anomalies and by assuming each anomaly to be due solely to an object five ft. tall, plus its shadow.

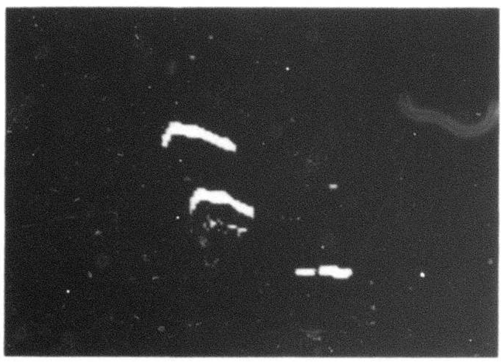

Figure 7 Shadows Found by Predictive Criterion

The third technique is based on a projective model. It tries to look directly for the shadow edge. Vehicles tend to be rectangular when viewed from above; and, unless the sun is directly ahead of or behind the vehicle, there will be a long, straight edge separating the vehicle

from its shadow. This edge can usually be found by performing a Hough transform [7] on the gradient of the image, or, equivalently, by projecting the gradient onto axes oriented in various directions and finding the direction from which the gradient points tend most to reinforce one another. However, much better results are obtainable when the direction of the edge is known or assumed a priori. Such is usually the case, for vehicles tend to be oriented parallel to the road direction.

An example of shadow detection by projection is presented in the next section.

The three techniques are based on different sets of assumptions and are applicable in different circumstances. The projective method is useful only for finding shadows of vehicles. The predictive model is more generally useful, being applicable to objects off the road as well as on it. The brightness model makes no assumptions about the object casting the shadow--it only requires that the background on which the shadow is cast be relatively uniform.

E. Classification of Anomalies

For classifying anomalies, we have chosen to construct a number of "expert" subroutines, each of which tests a specific hypothesis. For example, the vehicle expert determines whether or not a given anomaly could be a vehicle (plus its shadow) and if so, attempts to distinguish whether the vehicle is a car or a truck. The tree shadow expert tries to say whether or not the anomaly could be the shadow of an object off the road, and the road marking expert similarly looks for painted markings. Other expert modules could easily be integrated into the scheme. The experts operate in parallel, each expert forming its decision without interacting with its counterparts. The top-level program chooses the most likely interpretation of the anomaly. If no expert subroutine is able to account for the anomaly, it is labeled "unclassified."

The vehicle expert is the most involved of the expert subroutines. It first examines the overall size (area) of an anomaly. If the anomaly is too small or too large, it is rejected. Next, by projecting the gradient image to a baseline, long edges are found that might correspond to sides of the car. A binary mask is used for the projection, so that only those points near the anomaly are considered; the mask is generated by expanding ("growing") the anomaly region by three pixels. Figure 8a shows the results of applying a gradient operator to the image of Figure 5a. The masked gradient was projected on the axis drawn in Figure 8b, where the average projected gradient magnitude is plotted.

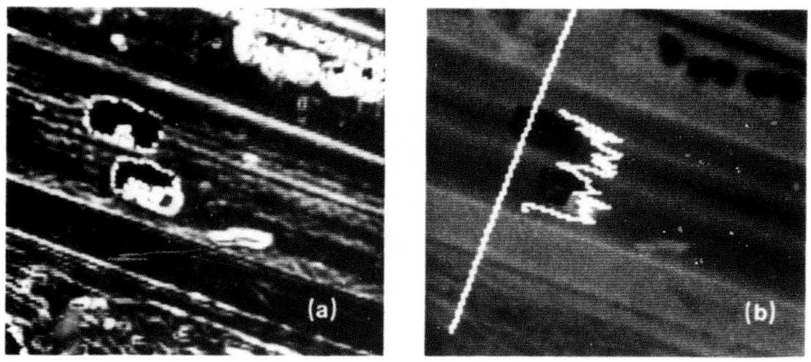

Figure 8 Use of Projection to Find Shadow Edges

A line perpendicular to the direction of the road is used as an initial baseline. If some evidence of edges is found, the orientation is perturbed a small amount to find a local maximum. If the edges are not found, a global search is made for a direction of projection that will show the edges. If the edges are again not found, the anomaly is rejected.

Note that there are three peaks in the plot, corresponding to the boundaries between road and car, between car and shadow, and between shadow and road. The three highest peaks in the projected gradient are examined to see if they are in the correct relationship. Average brightness is projected to the same baseline to see if the brightness of the shadow portion is appropriate. A figure of merit is computed from

these tests, indicating the degree to which the measured spacing and brightness approximate the expected spacing and brightness. The figure of merit is used later in choosing the most likely interpretation of the anomaly.

The average length of the shadow and the location of the sun may be used to estimate the height of the vehicle. A tolerance or range of uncertainty is also computed at this time, because the combination of low spatial resolution and a disadvantageous sun angle may make the height figure not particularly useful. A nominal height of 6 ft. is used for predicting a shadow to the front or the rear of the vehicle; this predicted shadow length subtracted from the length of the original anomaly yields the length of the vehicle.

Classification as to vehicle type is relatively crude at this time. If the overall length of the vehicle is greater than 20 ft., or if the height can be reliably stated as exceeding 6 ft., the vehicle is called a "truck." Otherwise it is called a "car."

Another expert subroutine identifies shadows of objects off the road. To qualify as such a shadow, an anomaly must have an average brightness lower than the average road brightness and extend to the edge of the road on the side nearer the sun. A figure of merit is calculated from the extent to which the average brightness (in the difference image) corresponds to the predicted value, as well as from the variance of brightness inside the anomaly.

The expert on painted road markings is similar to the shadow expert. Painted markings are always brighter than the road surface and limited in total area. The figure of merit is based only on variance of brightness; a much lower variance is expected for road markings than for shadows.

F. Discussion

The state of our experiments in anomaly classification is such that it is too early to report any quantitative results. However, we can say, qualitatively at least, that the methods outlined above succeed in the easy cases and break down for the difficult ones. We have tested our programs on approximately 20 different scenes extracted from three diverse road areas. Where good contrast exists between an anomaly and the road, and (in the case of vehicles) the shadow is visually distinct from the object casting it, we have little difficulty in obtaining a correct identification. Where conditions are not as good, the programs tend to make no identification at all, rather than come up with a misclassification. Additional robustness in the classifier will be necessary to enable it to handle unusual cases.

The various expert subroutines are not now integrated in any way. Each reports its figure of merit to the top-level program, which selects among the hypotheses. A more useful system should allow interaction among the various experts.

Figure 2 shows a good example of a case that could be handled by cooperation of the tree-shadow and the vehicle experts. It might be sufficient if the shadow expert were to realize that it could interpret part of the anomaly, subtract the explainable part, and ask the other experts to classify what remains. The vehicle expert would have to take the situation into account and not look for a separate shadow for this anomaly.

Figure 9 is difficult to analyze without higher-level knowledge. A more direct link to the data base would be particularly useful in this case, enabling us to divide the anomaly into portions that are "expected" (the visible portions of the arrow) and "not expected" (the car and its shadow).

Much generic knowledge tends to be expressed in the coding of the computer programs that analyze pictures. In this form it is inflexible--adding new knowledge involves writing new computer programs. A long-

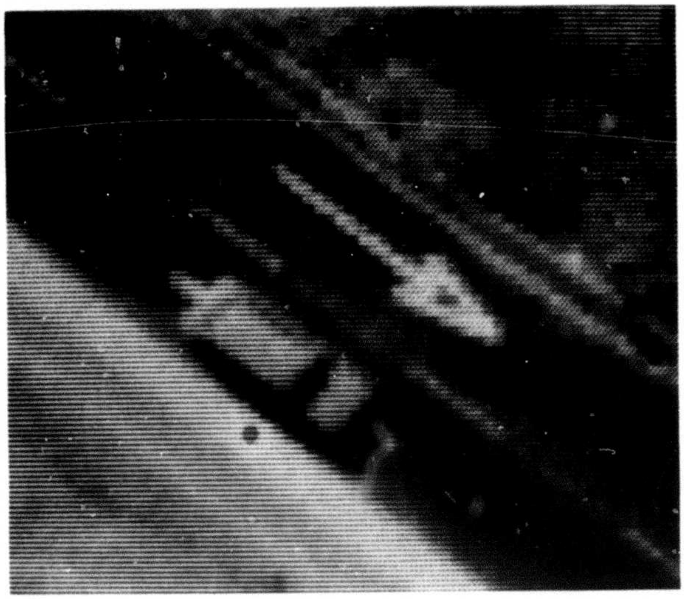

Figure 9 A Vehicle over a Road Marking

range goal of this research is to find new ways of expressing this kind of information--for example, in the form of rules or templates. Such a capability would lead to highly competent computer visual capabilities that would greatly enhance interactive and automatic cartography and photo interpretation.

REFERENCES

1. H. G. Barrow et al., "Interactive Aids for Cartography and Photo Interpretation: Progress Report, October 1977," Proceedings: Image Understanding Workshop, pp. 111-127 (October 1977).

2. M. A.Fischler et al., "Interactive Aids for Cartography and Photo Interpretation," Semiannual Technical Report, SRI Project 5300, SRI International, Menlo Park, California (October 1978).

3. L. H. Quam, "Road Tracking and Anomaly Detection," Proceedings: ARPA Image Understanding Workshop, pp. 51-55 (May 1978).

4. G. J. Agin, "Knowledge-Based Detection and Classification of Vehicles and Other Objects in Aerial Road Images," Proceedings: ARPA Image Understanding Workshop (in press, April 1979).

5. M. A. Fischler et al., "The SRI Image Understanding Program," Proceedings: ARPA Image Understanding Workshop, Menlo Park, California (April 1979).

6. R. C. Bolles et al., "The SRI Road Expert: Image-to-Data-Base Correspondence," Proceedings: ARPA Image Understanding Workshop, Pittsburgh, Pennsylvania (November 1978).

7. R. O. Duda and P. E. Hart, "Use of the Hough Transformation to Detect Lines and Curves in Pictures," Communications of the ACM, Vol. 15, No. 1 (January 1972).